PACKET RADIO TEMPORARY NOTE #117

Bob Metcalfe
Xerox PARC
October 16, 1974

WHO IS COUNTING PACKETS LOST DURING TURNAROUND?

That which we sheepishly call the Differential Turnaround Problem was recognized only long after our carrier-sense system was up and running. Because the Packet Radio Temporary Notes show no trace of turnaround considerations, we thought it best (though risky) to offer these friendly words of obvious advice so that you might not make the same miscalculations.

By "turnaround" we mean those actions taken in a station (or repeater) during the time from (1) when packet transmission, reception, or carrier sense is completed to (2) when the next of these operations can start. We must have incorrectly believed that the impact of real turnaround delays on performance would be negligible as long as such delays were kept small relative to packet duration or propagation delay. It is now clear that seemingly insignificant turnaround times can sometimes lead to substantial packet losses; the real surprise is that differences in turnaround time among operations are important.

Take the simple so-called 1-Persistent carrier-sense ALOHA scheme studied in numerous Packet Radio Temporary Notes. Say that only two terminals each were attempting to send a packet to a single destination station (or repeater). Assume that the carrier-sense hardware is working as intended and indeed one of the terminals is transmitting its packet while the other waits, deferring until carrier goes away. The transmission of the first packet comes to an end, its carrier goes away, and the second terminal, having detected absence of carrier, begins transmission. This would be carrier-sense in its finest hour were it not for turnaround or more specifically were it not for an "unfavorable" differential turnaround delay.

The destination's receiver has just now detected the end of the first packet; it quickly grabs its packet processor, checks for correct computed-on-the-fly "parity," drops status and length in memory with the preceding packet words, picks up the address pointer and word count of the next packet buffer, and turns on again; all in a relative instant, but too late. The second packet is already passing with its carrier pedestal erect, lost and desirous of retransmission.

The Differential Turnaround Problem is a result of a difference in complexity between two actions whose time differential must go the other way. In a straightforward implementation of a carrier-sense system by an otherwise clever computer engineer who is ignorant of the Differential Turnaround Problem, it will take less time (1) to sense that carrier has dropped and to start a waiting transmission than (2) to sense that carrier has dropped, to complete an incoming packet, and to start the reception of a new packet.

As suggested by the title of this note, the impact of turnaround losses in variously controlled multi-access systems has yet to be quantified. For now, just look at the unlikely but seemingly desirable case of packets arriving, one immediately after another, at a station (or repeater). The "unfavorable" but natural turnaround time differential just described would result in 50% packet losses.

Pending more careful analysis, we recommend the obvious. Terminals, repeaters, and stations should be constructed with a "favorable" turnaround time differential as follows. While unnatural from the strict digital-design standpoint, we should make it take longer for (1) a transmitter to sense loss of carrier and to begin packet transmission then for (2) a receiver to detect end of packet, to complete the incoming packet, and to start reception of the next. We waste a few microseconds at the end of each packet, but, by doing so, give receivers time to accept closely spaced packets.

It would seem that pure ALOHA and p-Persistent carrier-sense systems would be less susceptible to the above kind of Differential Turnaround Problem than the simple 1-Persistent or Busy Tone systems. If someone were to consider ignoring our advice and going ahead with the natural unfavorable turnaround differentials, then this susceptibility should be modelled.

A less important, even less believable, and yet real example of how we were surprised by differential turnaround will serve to add weight to our

Digitized by the Internet Archive
in 2025 with funding from
Amateur Radio Digital Communications, Grant 151

https://archive.org/details/temporary-note-117

argument for your attention. By building two kinds of station, one using modular programming in BCPL with limited hardware and the other using bummed assembly language with powerful hardware, we were able to stumble into a situation in which no data could be transferred from the modular BCPL station to the bummed assembly language station. We found that it took so long for the slow station to turn itself around to receive after sending a packet that it missed the quickly-returned acknowledgement from the fast station. We chose to speed up the send-receive turnaround of the slow station with some program bumming. Dick Binder reports that the ALOHA System's MENEHUNE keeps an acknowledgement turnaround delay table to solve a similar problem; in his case the slow send-receive turnaround in TCUs is due to their having a half-duplex transceiver in a two-frequency system.

Convinced there's something to this?

Dick Binder of The ALOHA System at the University of Hawaii (where I am now an appreciative guest) and Dave Boggs at Xerox PARC (where I live) have both discussed the subject of this note with me and are hereby relieved of responsibility should you think the note a total loss.